

Clinical Study: The Effect of Prednisone on Cholesterol

As a testbed for the prototype system we have been investigating the hypothesis that the steroid, prednisone, produces a significant elevation of plasma cholesterol. To test this hypothesis, the records of 50 patients with systemic lupus erythematosus (SLE) were transferred from the ARAMIS Database to SUMEX. Of these patients, 18 were found to have five or more cholesterol determinations and to have had sufficient variance in their prednisone regimens to be testable. The KB is used to elaborate a complex causal model for the prednisone/cholesterol hypothesis which is tested using a hierarchical multiple regression method with time-lagged values. The KB is used to determine sources of possible bias and to control for those variables in the regression or to eliminate corresponding time-intervals from records. An empirical Bayes method is used to average the estimated effects in patients with varying amounts of data.

The result, a highly statistically significant elevation of cholesterol by prednisone, will be submitted for publication during the coming year.

## 2. Research In Progress

Much work remains to be done in expanding the system software and in expanding the knowledge base. Current work is addressed to increasing the flexibility of the time-segmentation functions and enriching the data structures which encode relationships among objects.

We are trying to make increasingly general the class of medical hypotheses which the system can analyze automatically. This requires incorporating knowledge of additional statistical methods into the KB. We are also attempting to generalize our algorithms for selecting the set variables which may potentially confound a given hypothesis. As a means for testing and expanding the system's capabilities we intend to perform several specific studies of importance in the management of the rheumatic diseases. Our study of the effect of prednisone on cholesterol was mentioned above. Other studies now being planned include the effect of chronic aspirin ingestion on liver function in rheumatoid arthritis, the specific incidence of infectious complications of steroids as a function of dose and duration, and the utility of various autoantibodies in the prediction of flares of SLE as compared to the utility of other indicators.

Finally, we are developing a methodology for discovering hypotheses of interest in the database using a heuristically guided search of non-parametric lagged cross correlations.

## D. Publications

Blum, Robert L.: Discovery and Representation of Causal Relationships from a Time-Oriented Clinical Database: The RX Project. Stanford University Doctoral Dissertation (in press), 1981

Blum, Robert L.: Displaying Clinical Data from a Time-Oriented Database.  
(in press) Computers in Biology and Medicine, 1981

Wiederhold, Gio: Databases in Healthcare. Compendium Series on Technology in Healthcare, sponsored by the Healthcare Technology Center, Univ. of Missouri, Columbia, Mo., also available as Stanford CS Report 80-790

Blum, Robert L.: Automating the Study of Clinical Hypotheses on a Time-Oriented Data Base: The RX Project. Submitted for publication to MEDINFO80, Tokyo, Japan, Oct. 1980

Blum, Robert L., Wiederhold, Gio: Inferring Knowledge from Clinical Data Banks Utilizing Techniques from Artificial Intelligence. Proc. of The 2nd Annual Symp. on Computer Applications in Medical Care, pp. 303 to 307, IEEE, Washington, D.C., November 5-9, 1978

#### E. Funding Support Status

- 1) A Computer-Based System for Advising Physicians on Clinical Therapeutics  
Robert L. Blum, M.D.: Awardee  
Post-Doctoral Research Fellowship in Clinical Pharmacology  
Pharmaceutical Manufacturers' Association Foundation  
Total award: \$32,500 (direct)  
Term: July 1, 1978 to June 30, 1980
- 2) Integrating Medical Knowledge and Clinical Data Banks  
Robert L. Blum, M.D.: Principal Investigator  
National Library of Medicine, New Investigator Award  
Total award: \$90,000 (direct)  
Term: July 1, 1979 to June 30, 1982
- 3) Integrating Medical Knowledge and Clinical Data Banks  
Gio C. M. Wiederhold, Ph.D.: Principal Investigator  
National Center for Health Services Research, Small Grants  
Total award: \$35,000 (direct)  
Term: April 1, 1979 to March 31, 1981

## II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

### A. Collaborations

Since our project is relatively new, we do not yet have public versions of the programs. There is, however, a large sphere of collaboration which we expect in the future. Once the RX program is developed, we would anticipate collaboration with all of the ARAMIS project sites in the further development of a knowledge base pertaining to the chronic arthritides. The ARAMIS Project at SCIP is used by a number of institutions around the country via commercial leased lines to store and process their data. These institutions include the University of California School of Medicine, San Francisco and Los Angeles; The Phoenix Arthritis Center, Phoenix; The University of Cincinnati School of Medicine; The University of Pittsburgh School of Medicine; Kansas University; and The

University of Saskatchewan. All of the rheumatologists at these sites have closely collaborated with the development of ARAMIS, and their interest in and use of the RX project is anticipated. We hasten to mention that we do not expect SUMEX to support the active use of RX as an on-going service to this extensive network of arthritis centers, but we would like to be able to allow the national centers to participate in the development of the arthritis knowledge base and to test that knowledge base on their own clinical data banks.

#### B. Interactions with Other SUMEX-AIM Projects

Several of the concepts incorporated into the design of the RX Project have been inspired by other SUMEX-AIM Projects. The RX knowledge base is similar to the Units Package of the MOLGEN PROJECT. The production rule inference mechanism used by us is similar to that in the MYCIN Project.

Several programs developed by the MYCIN group are regularly used by RX. These include disk hash file facilities, text editing facilities, and miscellaneous LISP functions. Regular communication on programming details is facilitated by the on-line mail system.

#### C. Critique of Resource Management

The SUMEX KI-10 has been severely overloaded for at least a year. Working in LISP is impossible during the day and is even difficult at times which were formerly low utilization times. This has forced us to rely increasingly on other local computation facilities.

The SUMEX resource management, per se, has always been accessible and cooperative in trying to provide our project with adequate resources subject to prevailing constraints.

### III. RESEARCH PLANS

#### A. Project Goals and Plans

The overall goal of the RX Project is to develop a computerized medical information system capable of accurately extracting medical knowledge pertaining to the therapy and evolution of chronic diseases from a database consisting of a collection of stored patient records.

##### 1. Short-Term Goals

Goals for the year August, 1980 to July, 1981 have been detailed in section IC. above on research in progress. To summarize that section, our main short-term goal is to generalize and refine our methods for labeling and retrieving time-intervals or episodes from individual patient records and to generalize the class of hypotheses which the system is capable of analyzing. This requires further refinements in RX's algorithms for choosing and controlling for variables which may potentially confound an hypothesis of interest.

## 2. Long-Range Goals: August, 1981 to July, 1986

There are two inter-related long-range goals of the RX Project: 1) automatic discovery of knowledge in a large time-oriented database and 2) provision of assistance to a clinician who is interested in testing a specific hypothesis. These tasks overlap to the extent that some of the algorithms used for discovery are also used in the process of testing an hypothesis.

We hope to make these algorithms sufficiently robust that they will work over a broad range of hypotheses and over a broad spectrum of data distributions in the patient records.

### B. Justification for Continued Use of SUMEX

Computerized clinical data banks possess great potential as tools for assessing the efficacy of new diagnostic and therapeutic modalities, for monitoring the quality of health care delivery, and for support of basic medical research. Because of this potential, many clinical data banks have recently been developed throughout the United States. However, once the initial problems of data acquisition, storage, and retrieval have been dealt with, there remains a set of complex problems inherent in the task of accurately inferring medical knowledge from a collection of observations in patient records. These problems concern the complexity of disease and outcome definitions, the complexity of time relationships, potential biases in compared subsets, and missing and outlying data. The major problem of medical data banking is in the reliable inference of medical knowledge from primary observational data.

We see in the RX Project a method of solution to this problem through the utilization of knowledge engineering techniques from artificial intelligence. The RX Project, in providing this solution, will provide an important conceptual and technologic link to a large community of medical research groups involved in the treatment and study of the chronic arthritides throughout the United States and Canada, who are presently using the ARAMIS Data Bank through the SCIP facility via TELENET.

Beyond the arthritis centers which we have mentioned in this report, the TOD (Time-Oriented Data Base) User Group involves a broad range of university and community medical institutions involved in the treatment of cancer, stroke, cardiovascular disease, nephrologic disease, and others. Through the RX Project, the opportunity will be provided to foster national collaborations with these research groups and to provide a major arena in which to demonstrate the utility of artificial intelligence to clinical medicine.

### SUMEX as a Resource:

To discuss SUMEX as a resource for program development, one need only compare it to the environment provided by our other resource, the IBM 370/3033 installation at SCIP - the major computing resource at Stanford. Of the programs which we use daily on SUMEX -INTERLISP, MSG, TVEDIT, BBD, LINK - there is nothing even approaching equivalence on the 370, despite

its huge user community. These programs greatly facilitate communication with other researchers in the SUMEX community, documentation of our programs, and the rapid interactive development of the programs themselves. The development of a program involving extensive symbolic processing and as large and complex as RX at the SCIP facility, would require a staff many times as large as ours. The SUMEX environment greatly increases the productive potential of a research group such as ours to the point where a large project like RX becomes feasible.

#### Computation Resources Required by RX:

##### Disk Allocation:

RX requires the use of two large data files which need to be kept on-line: the patient database (DB) and the knowledge base (KB). In the course of testing a hypothesis several other files are used: inverted files, source files for statistical processing, LISP SYSOUT files, etc. Our current total disk allocation of 1500 pages for all RX group members has been just adequate. In the future, with anticipated expansions in numbers of patients and size of the KB, we intend to request an increase of our total allocation to 2000 pages.

##### C. Other Computational Resources

It is clear that the scope of potential application of the RX Project is large. Within the term of the SUMEX-AIM grant projected through July, 1986, we anticipate the involvement of several of the national ARAMIS collaborating institutions in developing and testing arthritis knowledge bases which reflect their own patient populations and therapeutic biases. The current SUMEX machine configuration will not be able to support this national interaction because the central processors of the KI-10 are already taxed to the limit. Ours is among the SUMEX groups which would greatly benefit by the addition of one or more PDP-10 compatible machines, which could provide support to our anticipated national user community. Another resource which would be highly desirable is a faster and more reliable means for transferring data interactively between SUMEX and the SCIP IBM 370. Our current method utilizes a 2400 baud line with transmission from SCIP to SUMEX only, and is fraught with a high error rate. The addition of a reliable local network facility would greatly facilitate our ability to transfer patient files from SCIP to SUMEX.

##### D. Recommendations for Resource Development

SUMEX is heavily loaded everyday and almost every evening. Program research is next to impossible during those periods. Program development would be greatly facilitated by the addition of any resources which lessened this loading: upgrading the current machine to a KL or adding core to decrease page swapping.

## II.A.2    National AIM Projects

The following group of projects is formally approved for access to the AIM aliquot of the SUMEX-AIM resource or the Rutgers-AIM resource. Their access is based on review by the AIM Advisory Group and approval by the AIM Executive Committee.

II.A.2.1 Acquisition of Cognitive Procedures (ACT)

## Acquisition of Cognitive Procedures (ACT)

Dr. John Anderson  
Carnegie-Mellon University

I. SUMMARY OF RESEARCH PROGRAM

## A. Project Rationale

To develop a production system that will serve as an interpreter of the active portion of an associative network. To model a range of cognitive tasks including memory tasks, inferential reasoning, language processing, and problem solving. To develop an induction system capable of acquiring cognitive procedures with a special emphasis on language acquisition and problem-solving skills.

## B. Medical Relevance and Collaboration

1. The ACT model is a general model of cognition. It provides a useful model of the development of and performance of the sorts of decision making that occur in medicine.
2. The ACT model also represents basic work in AI. It is in part an attempt to develop a self-organizing intelligent system. As such it is relevant to the goal of development of intelligent artificial aids in medicine.

We have been evolving a collaborative relationship with James Greeno and Allan Lesgold at the University of Pittsburgh. They are applying ACT to modeling the acquisition of reading and problem solving skills. We have made ACT a guest system within SUMEX. ACT is currently at the state where it can be shipped to other INTERLISP facilities. We have received a number of inquiries about the ACT system. ACT is a system in a continual state of development but we periodically freeze versions of ACT which we maintain and make available to the national AI community.

## C. Highlights of Research Progress

Our ACTF system is a production system that operates in a semantic network data base. Our learning work has been focused on ways of increasing the power of production systems for performing various tasks. One class of learning mechanisms concern what we call knowledge compilation. This involves automatic mechanisms for creating productions that directly perform behavior that formerly required interpretative processing of knowledge in the semantic network. These compilation mechanisms also model the process by which human experts develop special purpose procedures to deal with the different types of problems that occur in their domain of expertise.

Another class of learning mechanisms are concerned with tuning existing procedures so that they apply more appropriately. There are various mechanisms concerned with extending or generalizing the range of application of a procedure. In the past year we have been working at reducing these different generalization processes to a common partial matching process. In addition to generalization, tuning occurs in the ACTF system by means of discrimination and composition. Discrimination is a process for restricting the range of applicability of a production. Composition attempts to build macro-operators out of a series of productions.

The third direction of our learning work has been concerned with developing a flexible strength-based set of conflict resolution rules. Here we are concerned with modelling the gradual improvement seen in human cognitive skills and also providing the system with the resilience so that it can recover from noise and changes in environmental contingencies.

We have been applying this theory in detail to a simulation of how students acquire proof skills in geometry. We have a more or less thorough analysis of how students learn new postulates of geometry; we initially use these postulates in an interpretative fashion, integrating them with prior knowledge; how they compile special purpose procedures that directly apply this knowledge to proof generation; and how these procedures become tuned with practice. This application has provided strong evidence for most of the learning developments in the ACT system. It has also forced us to develop formalisms for how planning and problem-solving should be structured within a production-system framework.

#### D. List of Project Publications

- [1] Anderson, J.R. Language, Memory, and Thought. Hillsdale, N.J.: L. Erlbaum, Assoc., 1976.
- [2] Kline, P.J. & Anderson, J.R. The ACTE User's Manual, 1976.
- [3] Anderson, J.R., Kline, P. & Lewis, C. Language processing by production systems. In P. Carpenter and M. Just (Eds.). Cognitive Processes in Comprehension. L. Erlbaum Assoc., 1977.
- [4] Anderson, J.R. Induction of augmented transition networks. Cognitive Science, 1977, 125-157.
- [5] Anderson, J.R. & Kline, P. Design of a production system. Paper presented at the Workshop on Pattern-Directed Inference Systems, Hawaii, May 23-27, 1977.
- [6] Anderson, J.R. Computer simulation of a language acquisition system: A second report. In D. LaBerge and S.J. Samuels (Eds.). Perception and Comprehension. Hillsdale, N.J.: L. Erlbaum Assoc., 1978.
- [7] Anderson, J.R., Kline, P.J., & Beasley, C.M. A theory of the acquisition of cognitive skills. In G.H. Bower (Ed.). Learning and Motivation, Vol. 13. New York: Academic Press, 1979.



- [8] Anderson, J.R., Kline, P.J., & Beasley, C.M. Complex Learning. In R. Snow, P.A. Frederico, & W. Montague (Eds.). Aptitude, Learning, and Instruction: Cognitive Processes Analyses. Hillsdale, N.J.: Lawrence Erlbaum Assoc., 1980.
- [9] Anderson, J.R. & Kline, P.J. A learning system and its psychological implications. In the Proceedings of the Sixth International Joint Conference on Artificial Intelligence, 1979.
- [10] Reder, L.M. & Anderson, J.R. Use of thematic information to speed search of semantic nets. Proceedings of the Sixth International Joint Conference on Artificial Intelligence, 1979, 708-710.
- [11] Neves, D.M. & Anderson, J.R. Knowledge compilation: Mechanisms for the automatization of cognitive skills. In J.R. Anderson (Ed.), Cognitive Skills and their Acquisition. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1981.
- [12] Anderson, J.R., Greeno, J.G., Kline, P.J., & Neves, D.M. Acquisition of Problem-solving skill. In J. R. Anderson (Ed.), Cognitive Skills and their Acquisition. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1981.

#### E. Funding Support

An information-processing analysis of learning in geometry  
John R. Anderson, Principal Investigator  
National Science Foundation (IST-80-15357)  
\$186,000 Feb 15, 1981 - Feb 15, 1984

## II. INTERACTION WITH THE SUMEX-AIM RESOURCE

### A. Collaborations, Interactions, and Sharing of Programs via SUMEX.

We have received and answered many inquiries about the ACT system over the ARPANET. This involves sending documentations, papers, and copies of programs. The most extensive collaboration has been with Greeno and Lesgold who are also on SUMEX (see the report of the Simulation of Comprehension Processes project). There is an ongoing effort to assist them in their research. Feedback from their work is helping us with system design.

We find the SUMEX-AIM workshops (those that we could manage to attend) ideal vehicles for updating ourselves on the field and for getting to talk to colleagues about aspects of their work of importance to us.

Due to memory space problems encountered by ACT we expect that soon we will need to make use of the smaller version of INTERLISP developed at SUMEX for use in the CONGEN program.

### B. Critique of Resource Management

The SUMEX-AIM resource has been well suited for the needs of our project. We have made the most extensive use of the INTERLISP facilities and the facilities for communication on the ARPANET. We have found the SUMEX personnel extremely helpful both in terms of responding to our immediate emergencies and in providing advice helpful to the long-range progress of the project. Despite the fact that we are not located at Stanford, we have not encountered any serious difficulties in using the SUMEX system; in fact, there are real advantages in being in the Eastern time zone where we can take advantage of the low load on the system during the morning hours. We have been able to get a great deal of work done during these hours and try to save our computer-intensive work for this time.

Two location changes by the ACT project (from Michigan to Yale in the summer of 1976 and from Yale to Carnegie-Mellon in the summer of 1978) have demonstrated another advantage of working on SUMEX: In both cases we were back to work on SUMEX the day after our arrival.

## III. RESEARCH PLANS (8/80-7/86)

### A. Project Goals and Plans

Our long-range goals are: (1) Continued development of the ACT system; (2) Application of the system to modeling of various cognitive processes; (3) Dissemination of the ACT system to the national AI community.

This is a period of major evolution for the ACT theory. We have been developing three special versions of the ACTF learning that allow us to more efficiently simulate learning in three domains: proving theorems in geometry, speaking a new language, and writing programs in LISP. We are also performing special purpose simulations of the processes of spreading activation in memory retrieval and of pattern-matching processes in reading. We will be assimilating our experiences with these special purpose simulations in putting forth a major revision of the ACT theory. A research monograph is being written setting forth this theory and is scheduled for completion in late 1982. Subsequent to the writing of this monograph we intend to create an ACTG successor to ACTF that will embody the new conceptions.

### B. Justification for Continued Use of SUMEX:

Our goal for the ACT system is that it should serve as a ready-made "programming language" available to members of the cognitive science community for assembling psychologically-accurate simulations of a wide range of cognitive processes. Our intention and ability to provide such a resource justifies our use of the SUMEX facility. This facility is designed expressly for the purpose of developing and supporting such national AI resources and is, in this regard, clearly superior to the facilities we have available locally from the Carnegie-Mellon computer

science department. Among the most important SUMEX advantages are the availability of INTERLISP on a machine accessible by either the ARPANET or TYMNET and the existence of a GUEST login. It appears that, at least for the time being, ACT has no hope of being a national resource unless it resides at SUMEX and, given the local unavailability of a network-accessible INTERLISP, it would even be very difficult to shift any significant portion of our development work from SUMEX to CMU.

#### C. Needs and Plans for Other Computational Resources

Carnegie-Mellon's plans to begin upgrading its PDP-10 hardware to emerging state-of-the-art machines (VAX, LISP machines, etc.) promises to provide a excellent resource eventually, and we hope to have access to that resource as it develops. However, given that a considerable amount of software development will be required, a sophisticated LISP system such as INTERLISP is not likely to be available on this hardware in the near future.

#### D. Comments and Suggestions for Future Resource Goals

We are beginning to feel squeezed by various limitations of the SUMEX facility. The problem of peak load is quite serious. We have also been struggling with the address limitations of the current INTERLISP which is made more grievous by the amount of space INTERLISP requires. The computation time and address space limitations have meant that we have not been able to pursue certain projects that we would have otherwise. We applaud any efforts to increase computational power, to increase the address space of INTERLISP (e.g. VAXes), or to create significantly more space efficient versions of INTERLISP.

II.A.2.2 CADUCEUS Project (INTERNIST)

## CADUCEUS Project (\*)

J. D. Myers, M.D. and H. Pople, Ph.D.  
University of Pittsburgh  
Pittsburgh, Pennsylvania

I. SUMMARY OF RESEARCH PROGRAM

## A. Medical Rationale

The principal objective of this project is the development of a high-level computer diagnostic program in the broad field of internal medicine as an aid in the solution of complex and complicated diagnostic problems. To be effective, the program must be capable of multiple diagnoses (related or independent) in a given patient.

A major achievement of this research undertaking has been the design of a program called INTERNIST-I, along with an extensive medical data base now encompassing over 500 diseases and some 3450 individual manifestations of disease.

Although this consultative program is designed primarily to aid skilled internists in complicated medical problems, the program may have spin-off as a diagnostic and triage aid to physicians assistants, rural health clinics, military medicine and space travel.

Development of the INTERNIST-I system was begun about ten years ago. The system was successfully demonstrated for the first time in 1974 and has been used since that time in the analysis of hundreds of clinical problems.

A major point of departure for the design of the original INTERNIST program was the realization that the task of clinical decision making in internal medicine is an ill-structured problem. In other domains, the task of diagnosis is often viewed as one of pattern recognition or discrimination: there is available a predefined collection of possible classifications (characterizing disease entities or clinical states), one and only one of which is considered possible in the case being studied. A diagnostic problem solver dealing with such a well structured domain has the fairly straightforward task of selecting that one of this fixed set of alternatives which best fits the facts of the case. Many statistical,

-----  
(\*) For a variety of reasons, including a request from an agency alleging a prior claim on the name, future generations of the diagnostic program originally called INTERNIST will subsequently be referred to as CADUCEUS. This universal symbol of the medical profession seems appropriate to the expanded role we see for this type of program in the years to come. To avoid confusion in this report, the original program will continue to be called INTERNIST-I while references to the successor system, originally called INTERNIST-II, will now employ the new name.

pattern recognition, and algorithmic techniques have been employed successfully in performing computer aided diagnosis in these well structured clinical problem domains.

Primarily because complex cases often involve two or more concurrently active disease processes, no set of exhaustive and mutually exclusive classifications can be developed to structure the diagnostic problem in internal medicine. In principle, it might be argued that this more complex problem domain could be reduced to a simple discrimination task if, in addition to the individual disease entities, one includes appropriate multiple disease complexes in the set of allowable patient descriptors. However, since our experience indicates that as many as ten or twelve individual descriptors may apply in a complex clinical problem, and considering that there are a thousand or more individual descriptors of interest in Internal Medicine, the prospect of recording explicitly all possible multiple disease classifications is clearly infeasible.

Our thesis is that, in the absence of explicit structure derived from the problem domain, the successful clinician engages in heuristic imposition of structure so that effective problem solving strategies might be selected and employed for decision making relative to the postulated problem structure.

In INTERNIST-I, this concept of heuristic imposition of structure is expressed primarily by means of a novel "problem-formation" heuristic. In effect, the program composes dynamically, on the basis of evidence provided, what in context constitutes a presumed exhaustive and mutually exclusive subset of disease entities that can explain, more or less equally well, some significant subset of the observed findings in a clinical case. This heuristic problem structuring procedure is invoked repeatedly during the course of a diagnostic consultation in order to deal sequentially with the component parts of a complex clinical problem.

Because this program is intended to serve a consulting role in medical diagnosis, it has been challenged with a wide variety of difficult clinical problems: cases published in the medical journals, cpc's, and other interesting and unusual problems arising in the local teaching hospitals. In the great majority of these test cases, the problem-formation strategy of INTERNIST-I has proved to be effective in sorting out the pieces of the puzzle and coming to a correct diagnosis, involving in some cases as many as a dozen disease entities.

On the basis of this extensive test of the initial INTERNIST-I system, it has become clear that many aspects of the system's performance could be significantly enhanced if it would be possible to deal with the various component problems and their interrelationships simultaneously. This has led to the design of CADUCEUS, a system embodying strategies of concurrent problem-formation which we expect will yield more rapid convergence to the correct diagnosis in many cases, and in at least some cases provide more acceptable diagnostic behavior.

## B. Medical Relevance and Collaboration

The program inherently has direct and substantial medical relevance.

The institution of collaborative studies with other institutions has been deferred pending completion of the programs and knowledge base enhancements required for CADUCEUS. The installation of our own, dedicated VAX computer expected this summer will considerably aid future collaboration.

## C. Highlights of Research Progress

### Accomplishments This Past Year:

a) Prototypic computer programs have been written to operate CADUCEUS in the new diagnostic mode. The entire medical data base for the liver and biliary tract diseases has been reorganized into a form compatible with and utilizable by the CADUCEUS programs. Implementation of this work is pending the installation of the VAX computer when all of the programs must be written or rewritten using the FRANZ-LISP language.

b) The medical knowledge base comprising now just over 500 individual diseases and some 3450 manifestations of disease and hundreds of thousands of individual medical "facts," has been cumulative for the past eight years. Much effort has been spent during the past year in updating several dozens of diseases, most of which had been profiled years ago, and in establishing uniformity and consistency in this vast knowledge base. In addition, 17 new diseases have been profiled. The pediatric knowledge base has been expanded and now includes 78 diseases.

c) INTERNIST up to this time has been deficient in anatomic knowledge, particularly in topographical anatomy and anatomic laterality. An anatomic knowledge base beginning with neuroanatomy (the most complex) is being built for later incorporation into CADUCEUS. The knowledge base for the peripheral nervous system and the spinal cord is largely completed. The topographical anatomy of the abdomen and thorax are partially completed.

### Research in Progress:

There are five major components to the continuation of this research project:

- 1) The completion, continued updating, refinement and testing of the extensive medical knowledge base required for the operation of INTERNIST-I.
- 2) The completion and implementation of the improved diagnostic consulting program, CADUCEUS, which has been designed to overcome certain performance problems identified during the past five years' experience with the original INTERNIST-I program.

- 3) Institution of field trials of CADUCEUS on the clinical services in internal medicine at the Health Center of the University of Pittsburgh.
- 4) Expansion of the clinical field trials to other university health centers which have expressed interest in working with the system.
- 5) Adaptation of the diagnostic program and data base of CADUCEUS to subserve educational purposes and the evaluation of clinical performance and competence.

Current activity is devoted mainly to the first two of these, namely, the continued development of the medical knowledge base, and the implementation of the improved diagnostic consulting program (CADUCEUS). The development of the anatomic knowledge base is mentioned above.

Doctor Gordon Banks, a skilled neurologist who also has a Ph.D. in physics and considerable experience in computing, will be joining the team as of July 1, 1981 and will provide manpower and expertise for the further development of the sizeable and important neurological component of the medical knowledge base and its manipulation by the CADUCEUS programs.

#### D. List of Relevant Publications

1. Pople, H.E. "The Formation of Composite Hypotheses in Diagnostic Problem Solving: An Exercise in Synthetic Reasoning", Proceedings of the Fifth International Joint Conference on Artificial Intelligence, Boston, August 1977.
2. Pople, H.E. "On the Knowledge Acquisition Process in Applied A.I. Systems", Report of Panel on Applications of A.I., Proceedings of Fifth International Joint Conference on Artificial Intelligence, 1977.
3. Pople, H.E., Myers, J. D. & Miller, R.A. "The DIALOG Model of Diagnostic Logic and its Use in Internal Medicine, Proceedings of the Fourth International Joint Conference on Artificial Intelligence, Tbilisi, USSR, September 1975.
4. Pople, H.E. "Artificial Intelligence Approaches to Computer-Based Medical Consultation, Proceedings IEEE Intercon, New York, 1975.
5. Myers, J. D. "The Process of Clinical Diagnosis and Its Adaptation to the Computer," in The Logic of Discovery and Diagnosis in Medicine, University of Pittsburgh Series in the Philosophy and History of Science, University of California Press (in press).
6. Myers, J. D., Pople, H. E. & Miller, R. A. "INTERNIST: Can Artificial Intelligence Help?" in Clinical Decision and Laboratory Use, University of Minnesota Press (in press).

7. Pople, H. E. "Coming to Grips with the Multiple Diagnosis Problem," in Computer-Assisted Decision Making Using Clinical and Paraclinical (Laboratory) Data. B. Statland & S. Bauer (eds.) Mediad Inc., Tarrytown, N. Y., 1980, pp. 81-88. Reprinted in The Logic of Discovery and Diagnosis in Medicine, University of Pittsburgh Series in the Philosophy and History of Science, University of California Press (in press).
8. Pople, H. E. "Heuristic Methods for Imposing Structure on Ill-Structured Problems: The Structuring of Medical Diagnostics," in Artificial Intelligence in Medicine, AAAS Symposium Series, Westview Press (forthcoming 1981).

E. Funding support

1. Clinical Decision Systems Research Resource

Harry E. Pople, Jr., Ph.D.  
Associate Professor Business

Jack D. Myers, M.D.  
University Professor (Medicine)  
University of Pittsburgh

Division of Research Resources  
National Institutes of Health

2 R24 RR01101-04

07/01/80 - 06/30/85  
\$1,607,717

07/01/80 - 06/30/81  
\$465,199

2. INTERNIST: A Computer-Based Diagnostic Consultant

Harry E. Pople, Jr., Ph.D.  
Associate Professor of Business

Jack D. Myers, M.D.  
University Professor (Medicine)  
University of Pittsburgh

National Library of Medicine  
National Institutes of Health

1 R01 LM03710-01

07/01/80 - 06/30/85  
\$817,884

07/01/80 - 06/30/81  
\$148,458



### 3. New Computer-Based Patient Case Simulator

Randolph A. Miller, M.D.  
Associate Professor of Medicine  
University of Pittsburgh

National Library of Medicine - New Investigator  
National Institutes of Health

1 R23 LM03589-01

07/01/80 - 06/30/83  
\$89,350

07/01/80 - 06/30/81  
\$32,750

## II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

### A, B. Collaborations and Medical Use of Program Via SUMEX

CADUCEUS remains in a stage of research and development. As noted above, we are continuing to develop better computer programs to operate the diagnostic system, and the knowledge base cannot be used very effectively for collaborative purposes until it has reached a critical stage of completion. These factors have stifled collaboration via SUMEX up to this point and will continue to do so for the next year or two. In the meanwhile, through the SUMEX community there continues to be an exchange of information and states of progress. Such interactions particularly take place at the annual AIM Workshop.

### C. Critique of Resource Management

SUMEX has been an excellent resource for the development of CADUCEUS. Our large program is handled efficiently, effectively and accurately. The staff at SUMEX have been uniformly supportive, cooperative, and innovative in connection with our project's needs.

## III. RESEARCH PLANS (7/81-6/86)

### A. Project Goals and Plans

The prototype CADUCEUS programs and the trial reorganization of the liver and biliary tract diseases will be installed in the VAX over the summer and fall of this year. As rapidly as possible and pending further refinement and reorganization from experience with the new system, the remainder of the medical knowledge base will be entered. Local and later collaborative field trials must necessarily be postponed until this development has been accomplished.

At least 200 important medical diseases remain to be programmed. Renewed effort in this direction is now being expanded now that other tasks have been surmounted. Expanded efforts in the fields of neurology and pediatrics are included as described above.

#### B. Justification and Requirements for Continued SUMEX Use

Our use of SUMEX will obviously decline upon the installation of our VAX. Nevertheless, the excellent facilities of SUMEX are expected to be used for certain developmental work. It is intended, further, to keep INTERNIST-I at SUMEX for comparative use as CADUCEUS is developed here. Our team hopes to remain as a component of the SUMEX community and to share experiences and developments.

#### C. Needs and Plans for Other Computing Resources beyond SUMEX-AIM

Our predictable needs in this area will be met by the dedicated VAX computer soon to be installed.

#### D. Recommendations for Future Community and Resource Development

Whether a program like CADUCEUS, when mature, will be better operated from centralized, larger computers or from the developing self-contained personal computer is difficult to predict. For the foreseeable future it would seem that centralized, advanced facilities like SUMEX will be important in further program development and refinement.

II.A.2.3 Hierarchical Models of Human Cognition

## Hierarchical Models of Human Cognition (CLIPR Project)

Walter Kintsch and Peter G. Polson  
University of Colorado  
Boulder, Colorado

I. SUMMARY OF RESEARCH PROGRAM

## A. Project Rationale

The two CLIPR projects have made substantial progress in their research in this past year. This progress is almost completely due to our access to the SUMEX facility. The prose comprehension group has completed one major project, and is currently interacting with other SUMEX projects with the goal of building a prose comprehension model that reflects state-of-the-art knowledge from psychology and artificial intelligence.

The main activity of the planning group during the last year has been the detailed analysis of thinking-out-loud protocols collected from both expert and novice software designers. SUMEX facilities have been used to store, edit, and reformat the raw protocols to facilitate later analysis. Results of successive analyses are then input to SUMEX, and SUMEX facilities are used to collate the various results.

Technical Goals:

The CLIPR project consists of two subprojects. The first, the text comprehension project, is headed by Walter Kintsch and is a continuation of work on understanding of connected discourse that has been underway in Kintsch's laboratory for over seven years. The second, the planning project, is headed by Peter Polson of the University of Colorado and Michael Atwood of Science Applications Incorporated, Denver, and is studying the processes of planning using software design tasks.

The goal of the prose comprehension project is to develop a computer system capable of the meaningful processing of prose. This work has been generally guided by the prose comprehension model discussed by Kintsch and van Dijk (1978), although our programming efforts have identified necessary clarifications and modifications in that model (Miller & Kintsch, 1980a). Our more recent research (Miller & Kintsch, 1980b) has emphasized the importance of knowledge and knowledge-based processes in comprehension, and we are accordingly working with the AGE and UNITS groups at SUMEX toward the development of a knowledge-based, blackboard model of prose comprehension. We hope to be able to merge the substantial artificial intelligence research on these systems with psychological interpretations of prose comprehension, resulting in a computational model that is also psychologically respectable.

The primary goal of the planning project is the development of a model of human performance on software design tasks. We intend to begin by modeling protocols of experts on solving a particular problem, eventually extending the model to other levels of experience and problems. We propose a two-pronged attack on the process of developing a model.

The first is to develop a deeper understanding of our protocol data, to increase our knowledge of the details of the planning processes and the knowledge structures that experts use in the process of planning. We have developed a method of protocol analysis that essentially involves the transforming of the protocol into a low level theoretical description of the processes used to solve the design problem. We have assumed a very simplified version of a blackboard model that is described in Atwood and Jeffries (1980). We currently carry out our analysis by hand, developing a form of this low level model for each protocol. However, much of the activities involved in developing this model are clerical in nature and involve the categorization of segments of a verbal protocol and then the reorganization of the categorized information. Much of this work can be automated, and we propose to develop a program that will facilitate our protocol analysis and the development of the low level models that we use to describe the behavior of individual subjects.

Our second and much longer term objective is the development of a substantive model in AGE that can simulate the design processes. We feel that the software tools that are being developed at SUMEX -- in particular AGE and the UNITS package -- will dramatically facilitate our ability to develop this substantive model. Furthermore, current theoretical ideas about both the process of design and the representation of knowledge involved in developing a design have been strongly influenced by the MOLGEN project at SUMEX (Stefik, 1980).

#### B. Medical Relevance and Collaboration

The text comprehension project impacts indirectly on medicine, as the medical profession is no stranger to the problems of the information glut. By adding to the research on how computer systems might understand and summarize texts, and determining ways by which the readability of texts can be improved, medicine can only be helped by research on how people understand prose. Development of a more thorough understanding of the various processes responsible for different types of learning problems in children and the corresponding development of a successful remediation strategy would also be facilitated by an explicit theory of the normal comprehension process.

Note that our goal of a blackboard model is particularly relevant to the understanding of learning difficulties. One important aspect of a blackboard model is the separation of cognitive processes into a set of interacting subprocesses. Once such subprocesses have been identified and constructed, it would be instructive to observe the model's performance when certain of these processes are facilitated or inhibited. Many researchers have shown that there are a variety of cognitive deficits (insufficient short-term memory capacity, poor long-term memory retrieval, and such) that can lead to reading problems. Having a blackboard model in

which the power of individual components could be manipulated would be a significant step in determining the nature of such reading problems.

The planning project is attempting to gain understanding of the cognitive mechanisms involved in design and planning tasks. The knowledge gained in such research should be directly relevant to a better understanding of the processes involved in medical policy making and in the design of complex experiments. We are currently using the task of software design to describe the processes underlying more general planning mechanisms that are also used in a large number of task oriented environments like policy making.

Both the text comprehension project and the planning project involve the development of explicit models of complex cognitive processes; cognitive modelling is a stated goal of both SUMEX and research supported by NIMH.

The on-going development of the prose comprehension model would not be possible without our collaboration with the AGE and UNITS research groups. We look forward to a continued collaboration, with, we hope, mutually beneficial results. Several other psychologists have either used or shown an interest in using an early version of the prose comprehension model, including Alan Lesgold of SUMEX's SCP project, who is exporting the system to the LRDC vax. Needless to say, all of this interaction has been greatly facilitated by the local and network-wide communication systems supported by SUMEX. There has been considerable communication between members of the prose comprehension and AGE/UNITS groups as program bugs have been discovered and corrected; the presence of a mail system has made this process infinitely easier than if telephone or surface mail messages were required. The mail system, of course, has also enabled us to maintain professional contacts established at conferences and other meetings, and to share and discuss ideas with these contacts.

### C. Progress Summary

The prose comprehension project has completed an initial version of a model of prose comprehension (Miller & Kintsch, 1980a). This model has been applied to a large number of texts, and has yielded quite reasonable predictions of recall and readability. Psychologists from other universities have used this system to derive reading time and recall predictions for their own experimental materials; publication of this work is pending. We are currently using the AGE and UNITS packages to extend this model toward one that can make use of world knowledge in its analyses.

The planning group has completed the detailed analysis of several long thinking-out-loud protocols collected from both expert and novice software designers. These analyses involved the development of a lower level model for each of the protocols. See Atwood and Jeffries (1980) for details and examples.

## D. List of Relevant Publications

- Atwood, M. E., & Jeffries, R. Studies in plan construction I: Analysis of an extended protocol. Technical Report SAI-80-028-DEN, Science Applications, Incorporated, Denver, Co. March, 1980.
- Atwood, M. E., & Jeffries, R. Studies in plan construction II: Novice design behavior. Technical Report SAI-80-154-DEN, Science Applications, Incorporated, Denver, Colorado, December, 1980.
- Polson, P. G., Jeffries, R., Turner, A., & Atwood, M. E. The process of designing software. To appear in J. R. Anderson (Ed.), Cognitive skills and their acquisition. Hillsdale, N.J.: Erlbaum.
- Atwood, M. E., Polson, P. G., Jeffries, R., and Ramsey, H. R. Planning as a process of synthesis. Technical Report SAI-78-144-DEN, Science Applications, Incorporated, Denver, Co. December, 1978.
- Kintsch, W. On modelling comprehension. Invited address at the American Educational Research Association convention. San Francisco, April 10, 1979.
- Kintsch, W. and van Dijk, T. A. Toward a model of text comprehension and production. Psychological Review, 1978, 85, 363-394.
- Miller, J. R., & Kintsch, W. Readability and recall of short prose passages: A theoretical analysis. Journal of Experimental Psychology: Human Learning and Memory, 1980, in press.
- Miller, J. R., & Kintsch, W. Readability and recall of short prose passages. Text, 1981, in press.
- Miller, J. R. A knowledge-based mode of prose comprehension: Applications to expository text. Paper presented at the American Educational Research Association meeting, April, 1981.

## E. Funding Support Status

1. Readability and Comprehension.  
Walter Kintsch, Professor, University of Colorado  
National Institute of Education  
NIE-G-78-0172  
9/1/78 - 8/31/81: \$96,627  
9/1/80 - 8/31/81: \$46,537
2. Text Comprehension and Memory  
Walter Kintsch, Professor, University of Colorado  
National Institute of Mental Health  
5 Ro1 MH15872-9-13  
6/1/76 - 5/31/81: \$159,060  
6/1/80 - 5/31/81: \$32,880

3. Comprehension and Analysis of Information in Text  
Walter Kintsch, Professor, University of Colorado, and  
Lyle E. Bourne, Jr., Professor, University of Colorado  
Office of Naval Research, Personnel and Training Programs  
ONR N00014-78-C-0433  
6/1/78 - 5/31/80: \$68,315  
6/1/80 - 5/31/81: \$60,000
4. Procedural Net Theories of Human Planning and Problem Solving  
Michael Atwood, Research Psychologist, Science Applications,  
Incorporated, Denver, Colorado  
Office of Naval Research, Personnel and Training Programs  
ONR N0014-78-C-0165  
1/25/78 - 12/31/80: \$230,000  
1/1/80 - 6/30/81: \$85,000

## II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

### A. Sharing and Interactions with Other SUMEX-AIM Projects

Our primary interaction with the SUMEX community has been the work of the prose comprehension group with the AGE and UNITS projects at SUMEX. Feigenbaum and Nii have visited Colorado, and one of us (Miller) recently attended the AGE workshop at SUMEX. Both of these meetings have been very valuable in increasing our understanding of how our problems might best be solved by the various systems available at SUMEX. We also hope that our experiments with the AGE and UNITS packages have been helpful to the development of those projects.

We should also mention theoretical and experimental insights that we have received from Alan Lesgold and other members of the SUMEX SCP project. The initial comprehension model (Miller & Kintsch, 1980) has been used by Dr. Lesgold and other researchers at the University of Pittsburgh, as well as researchers at Carnegie-Mellon University, the University of Manitoba, Rockefeller University, and the University of Victoria.

### B. Critique of Resource Management

The SUMEX-AIM resource is clearly suitable for the current and future needs of our project. We have found the staff of SUMEX to be cooperative and effective in dealing with special requirements and in responding to our questions. The facilities for communication on the ARPANET have also facilitated collaborative work with investigators throughout the country.

## III. RESEARCH PLANS (8/79 - 7/81)

### A. Long Range Projects Goals and Plans

The primary long-term goal of the prose comprehension group is the development of a blackboard-based model of prose comprehension. Correspondingly, we anticipate continued use of the AGE and UNITS packages.

These packages allow us to model the knowledge structures possessed by people and the inferential processes that operate upon those structures, and are essential to our work.

The primary goal of the planning project is the development of a model, or a series of models, of human performance on the software design task. We intend to begin by modeling the protocols of experts on a particular task, eventually extending the model to other levels of experience and other tasks. To do this we will have to become more familiar with AGE and work on articulating our theory in a way that is compatible with the AGE framework. This will involve two parallel lines of effort. One is a deeper analysis of our protocol data, to increase our knowledge of the detailed planning processes and knowledge structures experts are using to solve these problems. The second is the development of a model in AGE that can simulate these processes. We have to date been using SUMEX only for the latter activity, but we are beginning to discover that both objectives are so intertwined that it is counter-productive for us to be using separate computer systems. We have transferred much of our protocol analyses activities to SUMEX, making it easier for us to share this very rich data source with other investigators.

#### B. Justification and Requirements for Continued SUMEX Use

The research of the prose comprehension project is clearly tied to continued access to the AGE and UNITS packages, which are simply not available elsewhere. We hope that our continued use of these systems will be offset by the input we have been and will continue to provide to those projects: our relationship has been symbiotic, and we look forward to its continuation.

#### C. Needs and Plans for Other Computational Resources

We currently use two other computing systems located at the University of Colorado. One is the Department of Psychology's VAX 11/780, which is used primarily to run real-time experiments to be modeled on SUMEX. The second is the University of Colorado's CDC 6400, which is used for various types of statistical analysis.

When the ARPA-sponsored Vax/Interlisp project is completed, we would be most interested in experimenting with becoming a remote AGE/UNITS site. It would seem that this sort of development is the ultimate goal of the package projects, and this type of interaction, once it becomes feasible, would be a logical extension of our association with the SUMEX facility.

#### D. Recommendations for Future Community and Resource Development

Our primary recommendation for future development within SUMEX involves (a) the continued support of INTERLISP, which is needed for AGE and for other work we have underway on SUMEX and (b) the continued development of the AGE and UNITS projects. In particular, we would like to see an extension of AGE to include a wider variety of control structures so that our psychological models would not be confined to one particular view of knowledge-based processing. The limited physical capacity of SUMEX,



both in terms of address space and overloading, is, as before, a major problem. The prose comprehension group can no longer use the publicly released AGE/UNITS system due to its severely limited address space, and has had to build a personal AGE system from a stripped-down version of Interlisp and a selected subset of AGE and UNITS. We heartily endorse the plans underway to obtain more computing capacity for the SUMEX project.

Given our acquisition of a VAX, we particularly support the ongoing and continued development of INTERLISP for the VAX, so that local use of AGE and UNITS would be possible. Since we, as well as other psychologists, need the real-time capability of VAX/VMS to run on-line experiments, we hope that the INTERLISP system to be developed will be compatible with VMS. Note that this need for real-time work coincides with real-world applications of SUMEX programs, in which a VAX might be devoted to both real-time patient monitoring and diagnostic systems such as PUFF or MYCIN.